arXiv:math/0211396v1 [math.GR] 26 Nov 2002

On the Properties of the Cayley Graph of Richard Thompson's Group F

V. S. Guba*

Vologda State Pedagogical University,
6 S. Orlov Street,
Vologda
Russia
160600
E-mail: guba@uni-vologda.ac.ru

February 1, 2008

Abstract

We study some properties of the Cayley graph of the R. Thompson's group F in generators x_0, x_1. We show that the density of this graph, that is, the least upper bound of the average vertex degree of its finite subgraphs is at least 3. It is known that a 2-generated group is not amenable if and only if the density of the corresponding Cayley graph is strictly less than 4. It is well known this is also equivalent to the existence of a doubling function on the Cayley graph. This means there exists a mapping from the set of vertices into itself such that for some constant $K > 0$, each vertex moves into the distance at most K and each vertex has at least two preimages. We show that the density of the Cayley graph of a 2-generated graph does not exceed 3 if and only if the group satisfies the same condition with $K = 1$. Besides, we give a very easy formula to find the length (norm) of a given element of F in generators x_0, x_1. This simplifies the algorithm by Fordham. The length formula may be useful to find the general growth function of F in generators x_0, x_1 and the growth rate of this function. In this paper we show that the lower bound for the growth rate of F is $(3 + \sqrt{5})/2$.

Introduction

The Richard Thompson group F can be defined by the following infinite group presentation

$$\langle\, x_0, x_1, x_2, \ldots \mid x_j x_i = x_i x_{j+1}\ (i < j)\ \rangle. \tag{1}$$

*This research is partially supported by the RFFI grant 99–01–00894 and the INTAS grant 99–1224.

This group was found by Richard J. Thompson in the 60s. We refer to the survey [6] for details. (See also [2, 3, 4].) It is easy to see that for any $n \geq 2$, one has $x_n = x_0^{-(n-1)} x_1 x_0^{n-1}$ so the group is generated by x_0, x_1. It can be given by the following presentation with two defining relations

$$\langle\, x_0, x_1 \mid x_1^{x_0^2} = x_1^{x_0 x_1}, x_1^{x_0^3} = x_1^{x_0^2 x_1} \,\rangle, \tag{2}$$

where $a^b = b^{-1}ab$ by definition.

Each element of F can be uniquely represented by a *normal form*, that is, an expression of the form

$$x_{i_1} x_{i_2} \cdots x_{i_s} x_{j_t}^{-1} \cdots x_{j_2}^{-1} x_{j_1}^{-1}, \tag{3}$$

where $s, t \geq 0$, $0 \leq i_1 \leq i_2 \leq \cdots \leq i_s$, $0 \leq j_1 \leq j_2 \leq \cdots \leq j_t$ and the following is true: if (3) contains both x_i and x_i^{-1} for some $i \geq 0$, then it also contains x_{i+1} or x_{i+1}^{-1} (in particular, $i_s \neq j_t$).

Equivalent definition of F can be done in the following way. Let us consider all strictly increasing continuous piecewise-linear functions from the closed unit interval onto itself. Take only those of them that are differentiable except at finitely many dyadic rational numbers and such that all slopes (derivatives) are integer powers of 2. These functions form a group under composition. This group is isomorphic to F. Another useful representation of F by piecewise-linear functions can be obtained if we replace $[0, 1]$ by $[0, \infty)$ in the previous definition and impose the restriction that all functions on infinity have the form $t \mapsto t + c$, where c is an integer.

The group F has no free subgroups of rank > 1. It is known that F is not elementary amenable (EA). However, the famous problem about amenability of F is still open. If F is amenable, then it is an example of a finitely presented amenable group, which is not EA. If it is not amenable, then this gives an example of a finitely presented group, which is not amenable and has no free subgroups of rank > 1. Note that the first example of a non-amenable group without free non-abelian subgroups has been constructed by Ol'shanskii [22]. (The question about such groups was formulated in [9], it is also often attributed to von Neumann [21].) Adian [1] proved that free Burnside groups with $m > 1$ generators of odd exponent $n \geq 665$ are not amenable. The first example of a finitely presented non-amenable group without free non-abelian subgroups has been recently constructed by Ol'shanskii and Sapir [23]. Grigorchuk [15] constructed the first example of a finitely presented amenable group not in EA.

The author thanks Matt Brin and Goulnara Arjantseva for helpful remarks.

1 Density

By the *density* of a finite graph Γ we mean the average value of the degree of a vertex in Γ. More precisely, let $v_1, \ldots, v_k$ be all of vertices of Γ. Let $\deg_\Gamma(v)$ denote the degree of a vertex v in the graph Γ, that is, the number of oriented edges of Γ that come out of v. Then

$$\delta(\Gamma) = \frac{\deg_\Gamma(v_1) + \cdots + \deg_\Gamma(v_k)}{k} \tag{4}$$

is the density of Γ.

Let G be a group generated by a finite set A. Let $C(G, A)$ be the corresponding (right) Cayley graph. Recall that the set of vertices of this graph is G and the set of edges is $G \times A^{\pm 1}$. For an edge $e = (g, a)$, its initial vertex is g, its terminal vertex is ga, and the inverse edge is $e^{-1} = (ga, a^{-1})$. The *label* of e equals a by definition. For the Cayley graph $C = C(G, A)$ we define the number

$$\bar{\delta}(C) = \sup_{\Gamma} \delta(\Gamma), \tag{5}$$

where Γ runs over all finite subgraphs of $C = C(G, A)$. So this number is the least upper bound of densities of all finite subgraphs of C. If C is finite, then it is obvious that $\delta(C) = \bar{\delta}(C)$. So we may call $\bar{\delta}(C)$ the *density of the Cayley graph* C.

Recall that a group G is called *amenable* whenever there exists a finitely additive normalized invariant mean on G, that is, a mapping $\mu\colon \mathcal{P}(G) \to [0, 1]$ such that $\mu(A \cup B) = \mu(A) + \mu(B)$ for any disjoint subsets $A, B \subseteq G$, $\mu(G) = 1$, and $\mu(Ag) = \mu(gA) = \mu(A)$ for any $A \subseteq G$, $g \in G$. One gets an equivalent definition of amenability if only one-sided invariance of the mean is assumed, say, the condition $\mu(Ag) = \mu(A)$ ($A \subseteq G$, $g \in G$). The proof can be found in [13].

The class of amenable groups includes all finite and abelian groups. It is invariant under taking subgroups, quotient groups, group extensions, and ascending unions of groups. The closure of the class of finite and abelian groups under these operations is the class EA of *elementary amenable* groups. A free group of rank > 1 is not amenable. There are many useful criteria for (non)amenability [11, 20, 14]. We need to mention the two properties of a finitely generated group G that are equivalent to non-amenability.

NA_1. *If G is generated by m elements and C is the corresponding Cayley graph, then the density of C has no maximum value, that is, $\bar{\delta}(C) < 2m$.*

Note that if NA_1 holds for at least one finite generating set, then the group is not amenable and so the same property holds for any finite generating set. For the proof of this property, we need to use the well-known *Følner condition* [11]. For our reasons it is convenient to formulate this condition as follows.

Let C be the Cayley graph of a group. By $\mathrm{dist}(u, v)$ we denote the distance between two vertices in C, that is, the length of a shortest path in C that connects vertices u, v. For any vertex v and a number r let $B_r(v)$ denote the ball of radius r around v, that is, the set of all vertices in C on the distance at most r from v. For any set Y of vertices, by $B_r(Y)$ we denote the r-neighbourhood of Y, that is, the union of all balls $B_r(v)$, where v runs over Y. By ∂Y we denote the *boundary* of Y, that is, the set $B_1(Y) \setminus Y$. The Følner condition (for the case of a finitely generated group) says that G is amenable whenever $\inf \#\partial Y / \#Y = 0$, where the infimum is taken over all non-empty finite subsets in G in a Cayley graph of G in finite number of generators (this property does not depend on the choice of a finite generating set). Any finite set Y of vertices in C defines a finite subgraph (also denoted by Y). The degree of any vertex v in C equals $2m$, where m is the number of generators. We know that exactly $\deg_Y(v)$ of the $2m$ edges that come out of v, connect the vertex v to a vertex from Y. The other $2m - \deg_Y(v)$ edges connect v to a vertex from ∂Y. Note that

each vertex of ∂Y is connected by an edge to at least one vertex in Y. This implies that the cardinality of ∂Y does not exceed the sum $\sum(2m - \deg_Y(v))$ over all vertices of Y. Dividing by $\#Y$ (the number of vertices in Y) implies the inequality $\#\partial Y/\#Y \leq 2m - \delta(Y)$. If $\bar{\delta}(C) = 2m$, then Y can be chosen such that $\delta(Y)$ is arbitrarily close to $2m$ so $\#\partial Y/\#Y$ will be arbitrarily close to 0. On the other hand, for any vertex v in Y there are at most $2m$ edges that connect v to a vertex in Y. Therefore, the sum $\sum(2m - d_Y(v))$ does not exceed $2m\#\partial Y$. So $2m - \delta(Y) \leq 2m\#\partial Y/\#Y$. If the right hand side can be done arbitrarily close to 0, then $\delta(Y)$ approaches $2m$ so $\bar{\delta}(C) = 2m$.

NA_2. *If C is the Cayley graph of G in a finite set of generators, then there exists a function $\phi\colon G \to G$ such that a) for all $g \in G$ the distance $\mathrm{dist}(g, \phi(g))$ is bounded by a constant $K > 0$, b) any element $g \in G$ has at least two preimages under ϕ.*

An elegant proof of this criterion based on the Hall-Rado theorem can be found in [7], see also [10]. Note that this property also does not depend on the choice of a finite generated set. A function ϕ from NA_2 will be called a *doubling function* on the Cayley graph C.

We need a definition. Suppose that NA_2 holds for the Cayley graph of a group G for the case $K = 1$. Then we say that the Cayley graph C is *strongly non-amenable*. The function $\phi\colon G \to G$ will be called a *strong doubling function* on the Cayley graph C. Note that each vertex is either invariant under ϕ or it maps into a neighbour vertex. We know that NA_2 holds if and only if the group is not amenable, that is, $\bar{\delta}(C) < 2m$. Now we would like to find out what happens if the Cayley graph of a 2-generated group is strongly non-amenable.

Theorem 1 *The Cayley graph of a group with two generators is strongly non-amenable if and only if the density of this graph does not exceed* 3.

Proof. Let G be a group with 2 generators and let C be the corresponding Cayley graph. It follows from the proof of [7, Theorem 32] that C is strongly non-amenable if and only if the doubling inequality holds, that is, $\#B_1(Y) \geq 2\#Y$ for any finite set Y of vertices in C. Indeed, if C admits a doubling function ϕ from NA_2 with $K = 1$, then $\phi^{-1}(Y)$ has at least $2\#Y$ elements and it is contained in $B_1(Y)$. To prove the converse, one needs to consider a bipartite graph with two classes of vertices both equal to G. Two vertices from the different classes are connected by an edge whenever the distance between them in C does not exceed $K = 1$. The doubling inequality implies that the conditions of the Hall-Rado theorem hold. Therefore, the bipartite graph has a perfect $(2, 1)$-matching. This means that C admits a doubling function with $K = 1$.

The "only if" part is trivial since the doubling inequality is equivalent to the fact that $\#\partial Y \geq \#Y$ for any finite subset Y in C. We know that $\#\partial Y/\#Y \leq 2m - \delta(Y)$ for any non-empty finite subset $Y \subseteq C$, where $m = 2$. This means that $\delta(Y) \leq 2m - 1 = 3$ and so $\bar{\delta}(C) \leq 3$.

Now suppose that $\bar{\delta}(C) \leq 3$. Let Y be a finite subgraph of C with k vertices. For each $0 \leq s \leq 4$, let q_s be the number of vertices in Y that have degree s in Y. Clearly, $k = q_0 + q_1 + q_2 + q_3 + q_4$. The number of oriented edges in Y is the sum of all degrees of vertices, that is, $q_1 + 2q_2 + 3q_3 + 4q_4$. If we divide this number by k, then we get $\delta(Y)$. Since $\bar{\delta}(C) \leq 3$, we have $\delta(Y) \leq 3$, which is equivalent to the inequality $q_1 + 2q_2 + 3q_3 + 4q_4 \leq$

$3k = 3(q_0 + q_1 + q_2 + q_3 + q_4)$ and so it can be rewritten as $3q_0 + 2q_1 + q_2 - q_4 \geq 0$. For any finite subgraph Y, we denote the number $3q_0 + 2q_1 + q_2 - q_4$ by $q(Y)$.

Suppose that the doubling inequality does not hold. Let us choose a minimal counterexample Y to the doubling inequality, that is, a finite subgraph Y in C with the property $\#B_1(Y) < 2\#Y$ for which the number $q(Y) \geq 0$ takes the smallest possible value. Each vertex in $\partial Y = B_1(Y) \setminus Y$ is connected by an edge with at least one vertex in Y. Suppose that some vertex v in ∂Y is connected with at least two vertices in Y. Let Y' be the subgraph in C with the new vertex v added to Y and two new non-oriented edges that connect v with vertices v_1, v_2 in Y. By definition, $B_1(Y') = B_1(Y) \cup B_1(v)$. The vertex v has exactly 4 edges that come out of v. At least two of them connect v with a vertex in Y. So $B_1(v)$ may contain at most 2 vertices not in $B_1(Y)$ (the vertex v itself belongs to $B_1(Y)$). This means that $\#B_1(Y') \leq \#B_1(Y) + 2 < 2\#Y + 2 = 2\#Y'$. Hence Y' is also a counterexample. To complete the proof, we need to check that $q(Y') < q(Y)$.

When we add a non-oriented edge that connects v and v_j $(j = 1, 2)$, then the degree of v_j increases by 1. This means that if $\deg_Y(v_j) = i$, then $\deg_{Y'}(v_j) = i + 1$. So $0 \leq i < 4$, the value of q_i decreases by 1, the value of q_{i+1} increases by 1. Clearly, $q(Y) = 3q_0 + 2q_1 + q_2 - q_4$ decreases by 1. So if we add the edges for both v_1 and v_2 (this does not exclude the case $v_1 = v_2$), then $q(Y)$ decreases by 2. But we also have a new vertex v that has degree 2 in Y'. Thus q_2 increases by 1 and so after all these operations we have $q(Y') = q(Y) - 1$. This contradicts the minimality of Y.

The proof is complete.

Note that the proof of Theorem 1 goes without any changes if we apply it to any regular graph of degree 4. Also we have to mention that the density of a Cayley graph of a group is closely related to an isoperimetric constant ι_* of a graph (see the definition in [7]). Namely, one has the equality $\iota_*(C) + \bar{\delta}(C) = 2m$ for the Cayley graph C of an m-generated group.

Theorem 1 applied to the Cayley graph $\mathcal{C}_2$ of F in generators x_0, x_1 means that if we cannot find a subgraph in $\mathcal{C}_2$ with density greater than 3, then there exists a doubling function on $\mathcal{C}_2$. One can imagine this doubling function in the following way. Suppose that a bug lives in each vertex of $\mathcal{C}_2$. We allow these bugs to jump at the same time such that each bug either returns to its initial position or it jumps to a neighbour vertex. As a result, we must have at least two bugs in each vertex.

It is natural to ask how much the value of $\delta(Y)$ can be for the finite subgraphs we are able to construct. It is easy to see that in each finite subgraph Y of $\mathcal{C}_2$ there exists a vertex of degree at most 2. Indeed, if all vertices were of degree 3 or 4, then one could travel along Y by positively labelled edges only (if we enter a vertex by an edge labelled by x_0 or x_1, then we can leave this vertex travelling along an edge with one of these labels). But F does not have nontrivial relations that involve positive letters only.

However, the following result shows that we are able to construct finite subgraphs in $\mathcal{C}_2$ with the density arbitrarily close to 3.

Theorem 2 *Let $\mathcal{C}_2$ be the Cayley graph of F in generators x_0, x_1. For any integers $m, n \geq 1$ there exists a finite subgraph $\bar{\Gamma}_{n,m}$ in $\mathcal{C}_2$ such that*

$$\lim_{m\to\infty} \delta(\bar{\Gamma}_{n,m}) = \frac{6(n-1)}{2n-1} \tag{6}$$

for any $n \geq 2$. In particular, $\bar{\delta}(\mathcal{C}_2) \geq 3$.

Proof. It is convenient to work with subgraphs in $\mathcal{C} = \mathcal{C}_\infty$. Note that all Cayley graphs $\mathcal{C}_i$ $(i = 2, 3, \ldots)$ are subgraphs in $\mathcal{C}$. Let Y be a finite connected subgraph in $\mathcal{C}$. Assume that Y is *full*, that is, if two vertices of Y can be connected in $\mathcal{C}$ by an edge, then this edge belongs to Y. It is not hard to see that any finite subset in F determines the corresponding full subgraph, which will be finite.

Let Y be a finite connected full subgraph in $\mathcal{C}$. By a *rank* of a vertex v in Y we mean the maximum number $k \geq 0$ such that v has an edge in Y labelled by $x_k^{\pm 1}$ that comes out of v. Suppose that for some $i \geq 0$, the rank of any vertex in Y is greater than i. Under these conditions, we define a subgraph in $\mathcal{C}$ denoted by $\mathcal{A}_i(Y)$.

Let $v_1, \ldots, v_k$ be the list of all vertices in Y (they are elements of F). Let $r_1, \ldots, r_k$ be the ranks of these vertices (each of these numbers exceeds i). For any $1 \leq j \leq k$, we consider the set of vertices of the form $v_j x_i^{-s}$, where $0 \leq s < r_j - i$. We call this set a *column* of v_j. The number of vertices in this column equals $r_j - i$ (if $r - j = i + 1$, then the column consists of v_j only). The vertices in a column are connected in $\mathcal{C}$ by edges with label x_i. The number of these edges in the column of v_j equals $r_j - i - 1$. Note that all the columns are disjoint. The union of them will be the set of vertices of the graph $\mathcal{A}_i(Y)$ we want to define.

Now let us take an arbitrary edge from Y. Suppose that it connects v' and v''. Let r', r'' be the ranks of these vertices, respectively. Without loss of generality we may assume that our edge has label x_k, where $k > i$ so we have $v'x_k = v''$. If $k = i + 1$, then we do nothing with this edge. If $k = i + 2$, then $v'x_i^{-1}$ and $v''x_i^{-1}$ both appear in the columns of v' and v''. In the graph $\mathcal{C}$, these vertices are connected by an edge labelled by x_{i+1} since $v'x_i^{-1}x_{i+1} = v'x_{i+2}x_i^{-1} = v''x_i^{-1}$. In general, if $k = i + p + 1$, where $p \geq 1$, then the rank of each of v', v'' is at least $i + p + 1$. So the elements $v'x_i^{-t}$, $v''x_i^{-t}$ will appear in the columns for all $1 \leq t \leq p$. From the defining relations of F it follows that $v'x_i^{-t}$, $v''x_i^{-t}$ are connected in $\mathcal{C}$ by an edge with label x_{k-t}. Indeed, $v'x_i^{-t}x_{k-t} = v'x_k x_i^{-t} = v''x_i^{-t}$ (we used the fact $k - t > i$ and the relations of the form $x_i^{-1}x_j = x_{j+1}x_i^{-1}$ that hold in F for all $j > i$).

So for any edge e labelled by x_k from v' to v'' in Y $(k \geq i + 1)$, we have $k - i$ edges (including e) that connect a vertex in the column of v' with a vertex in the column of v''. These edges are called *parallel* to e. The union of all these edges over all edges e in Y forms the set of edges of $\mathcal{A}_i(Y)$.

So we have defined the new graph $\mathcal{A}_i(Y)$. It is finite, connected, and it contains Y. Let us check that $\mathcal{A}_i(Y)$ is full. Suppose that two vertices $v'x_i^{-s}$ and $v''x_i^{-t}$ of $\mathcal{A}_i(Y)$ are connected by an edge f with label x_k. We have inequalities $0 \leq s < r' - i$, $0 \leq t < r'' - i$, where r', r'' are the ranks of v', v'', respectively. The graph Y is connected so there exists a path from v' to v'' in Y. The label of this path is a word W that involves only letters of the form $x_j^{\pm 1}$, where $j > i$. Equalities $v'W = v''$ and $v'x_i^{-s}x_k = v''x_i^{-t}$ imply $x_i^s W = x_k x_i^t$. We consider several cases.

1) $k < i$. In this case x_k belongs to the subgroup of F generated by x_{k+1}, x_{k+2}, $\ldots$, which is impossible.

2) $k = i$. In this case $x_i^{t-s+1} = W$ belongs to the subgroup generated by x_{i+1}, x_{i+2}, $\ldots$. This can happen only if the exponent $t - s + 1$ equals 0 and $W = 1$ in F. So $v' = v'' = v$ and our vertices have the form vx_i^{-s}, $vx_i^{-(s+1)}$. These vertices are connected by the edge with label x_i in $\mathcal{A}_i(Y)$.

3) $k > i$. The word $W^{-1}x_i^{-s}x_kx_i^t$ equals 1 in F so the algebraic exponent sum over all letters with the smallest subscript must be zero. In our case, the smallest subscript is i so $s = t$. Therefore, $W = x_i^{-s}x_kx_i^s = x_{k+s}$ in F. The graph Y is full so the edge e labelled by x_{k+s} that connects vertices v', v'' must belong to Y. Then our edge f labelled by x_k is parallel to e and thus belongs to $\mathcal{A}_i(Y)$ by definition.

Note that the rank of each vertex in $\mathcal{A}_i(Y)$ is at least i. This means that we can apply the operator $\mathcal{A}_i'$ to $\mathcal{A}_i(Y)$ for any $0 \le i' < i$. We start with the "linear" graph $\Xi_{n,m}$ that consists of $m+1$ vertices 1, x_n^{-1}, ..., x_n^{-m} and m positive edges labelled by x_n that connect these vertices. It is obvious that $\Xi_{n,m}$ is full and all vertex ranks are equal to n. We can apply the operator $\mathcal{A}_{n-2}$ to it ($\mathcal{A}_{n-1}$ can be also applied but this is useless). After that, we can apply $\mathcal{A}_{n-3}$, ..., $\mathcal{A}_1$, $\mathcal{A}_0$. As a result we get the family of subgraphs

$$\Gamma_{n,m} = \mathcal{A}_0\mathcal{A}_1 \cdots \mathcal{A}_{n-2}\Xi_{n,m}. \tag{7}$$

Each of them is a subgraph in $\mathcal{C}_\infty$. If we erase all edges of $\Gamma_{n,m}$ that have labels of the form $x_k^{\pm 1}$, where $k > 1$, then we get the subgraph $\bar{\Gamma}_{n,m}$ in $\mathcal{C}_2$ from the statement of our Theorem.

It is possible to characterize all vertices of $\Gamma_{n,m}$ as some words with negative exponents. We only mention that all of them will be of the form $x_n^{-s_n}x_{n-2}^{-s_{n-2}} \cdots x_1^{-s_1}x_0^{-s_0}$, where $s_0, s_1, \ldots, s_{n-2}, s_n \ge 0$.

Let Z be a subgraph in $\mathcal{C}_\infty$. By a *star* of a vertex v in Z (or a *Z-star*) we mean the set of labels of all edges that come out of v and belong to Z. Suppose that Z', Z'' are two subsets of the set of vertices of Z and $\lambda\colon Z' \to Z''$ is a bijection. If the Z-star of each vertex $v \in Z'$ coincides with the Z-star of its image $\lambda(v)$, then we say that λ is a *local isomorphism* (in Z). Now let Y be a finite connected full subgraph in $\mathcal{C}_\infty$ such that all its vertices have rank $> i$. For any subset Y' of Y we can denote by $\mathcal{A}_i(Y')$ the union of columns of all vertices from Y'. (Note that each full subgraph is determined uniquely by the set of its vertices so this notation does not lead to a confusion. For a single vertex v we shall write $\mathcal{A}_i(v)$ instead of $\mathcal{A}_i(\{v\})$.) It is worth noting that if two vertices $v, w \in Y$ have the same Y-star, then they have the same rank r and there exists a natural bijection between their columns, $vx_i^{-k} \mapsto wx_i^{-k}$, where $0 \le k < r$. Obviously, the $\mathcal{A}_i(Y)$-stars of vx_i^{-k} and wx_i^{-k} are also the same. So we can conclude that if Y' and Y'' are disjoint subsets and λ is a local isomorphism between Y' and Y'' in Y, then the sets $\mathcal{A}_i(Y')$ and $\mathcal{A}_i(Y'')$ will be also disjoint and there exists a local isomorphism between them in $\mathcal{A}_i$.

Now we can consider two vertices x_n^{-s}, x_n^{-t} in $\Xi_{n,m}$, where $0 < s < t < m$. Obviously, the sets $\{x_n^{-s}\}$ and $\{x_n^{-t}\}$ are locally isomorphic disjoint subsets in $\Xi_{n,m}$. Applying the argument from the above paragraph, we see that the sets V_s and V_t are disjoint locally isomorphic subsets in $\Gamma_{n,m}$, where $V_k = \mathcal{A}_0\mathcal{A}_1 \cdots \mathcal{A}_{n-2}(\Xi_{n,m})$ by definition for any $0 \le k \le m$. It is clear that all the sets V_0, V_1, ..., V_m form the disjoint subdivision of the set of vertices of $\Gamma_{n,m}$. These sets have exactly the same number of vertices. Let ρ_k ($0 \le k \le m$) be the average degree of a vertex in the subset V_k, that is, the number $\sum_{v \in V_k} \deg(v)/\#V_k$ (the degree is taken in the whole Cayley graph). We have $\rho_1 = \cdots = \rho_{m-1} = \rho$. The density of $\Gamma_{n,m}$ will be equal to

$$\bar{\delta}(\Gamma_{n,m}) = \frac{\rho_0 + \rho_1 + \cdots + \rho_{m-1} + \rho_m}{m+1}. \tag{8}$$

Since ρ_0 and ρ_m are bounded, the limit of (8) as m approaches ∞, is exactly ρ. Thus to complete the proof, we need to show that $\rho = 6(n-1)/(2n-1)$.

To calculate ρ, it is convenient to introduce the new family Γ_n of auxiliary finite graphs (they are no longer subgraphs of $\mathcal{C}_\infty$). But they will be still labelled graphs (automata) and it will be clear from the definition of them that the set of vertices of Γ_n is in a bijection with the set of vertices of V_k for any $1 \le k < m$. Moreover, each vertex in Γ_n have the same star as the corresponding vertex in V_k. Thus ρ will be equal to $\bar{\delta}(\Gamma_n)$. The idea is to extend the notion of $\mathcal{A}_i$. Let Y be any labelled finite graph, where all edges have labels of the form $x_j^{\pm 1}$, $j \ge 0$. The rank of a vertex is defined in the same way. Suppose that all vertices in Y have rank $> i$. For each vertex v of rank r we consider the set of $r-i$ vertices v_0, v_1, …, v_{r-i-1}, where $v_0 = v$. These vertices form the column of v. We connect them by $r-i-1$ directed edges labelled by x_i, where the kth edge ($1 \le k < r-i$) goes from v_k to v_{k-1}. (It is easy to see that v_k here is an analog of vx_i^{-k} in the above construction.) Then for any edge e of Y that has label x_j and goes from v to w, we consider $j-i$ edges e_0, e_1, …, e_{j-i-1}. Namely, for any $0 \le k < j-i$, we connect v_k and w_k by an edge labelled by x_{j-k}. Clearly, $e_0 = e$. So we get a graph that contains Y and we denote it by $\mathcal{A}_i(Y)$, as above. It is easy to see that we have an extension of the above concept. Also it is clear that we can repeat applications of the operators $\mathcal{A}_i$ with decreasing subscripts. If we start with the graph Ξ_n that has a single vertex and a loop labelled by x_n at this vertex, then we get the graph

$$\Gamma_n = \mathcal{A}_0\mathcal{A}_1 \cdots \mathcal{A}_{n-2}\Xi_n.$$

By $\bar{\Gamma}_n$ we denote the graph obtain from Γ_n by erasing all edges labelled by $x_i^{\pm 1}$, where $i \ge 2$. Since the vertex of Ξ_n has the same star as the vertex x_n^{-k} in $\Xi_{n,m}$ for any $1 \le k < m$, we easily conclude from our definitions that $\delta(\bar{\Gamma}_n) = \rho_k = \rho$.

The graphs Γ_n ($n \ge 1$) can be easily drawn explicitly. If a vertex v has a loop at v labelled by x_m, then it will be clear that the rank of v equals m. So we can draw this vertex as a circle with the number m inside. If Y is a labelled graph with labels of the form x_j ($j \ge 0$), then by $\Psi(Y)$ we denote the graph obtained from Y by increasing all subscripts of the labels by 1. We know that Γ_1 is a single loop labelled by x_1. To obtain Γ_{n+1} from Γ_n ($n \ge 1$), one has to apply Ψ to Γ_n and then apply $\mathcal{A}_0$. This is true because $\Gamma_{n+1} = \mathcal{A}_0(\mathcal{A}_1 \cdots \mathcal{A}_{n-1}\Xi_{n+1}) = \mathcal{A}_0(\Psi(\mathcal{A}_0 \cdots \mathcal{A}_{n-2}\Xi_{n-1})) = \mathcal{A}_0(\Psi(\Gamma_n))$. This gives an easy way to imagine how these graphs look like. We illustrate this process by the following picture that shows Γ_n for $1 \le n \le 4$.

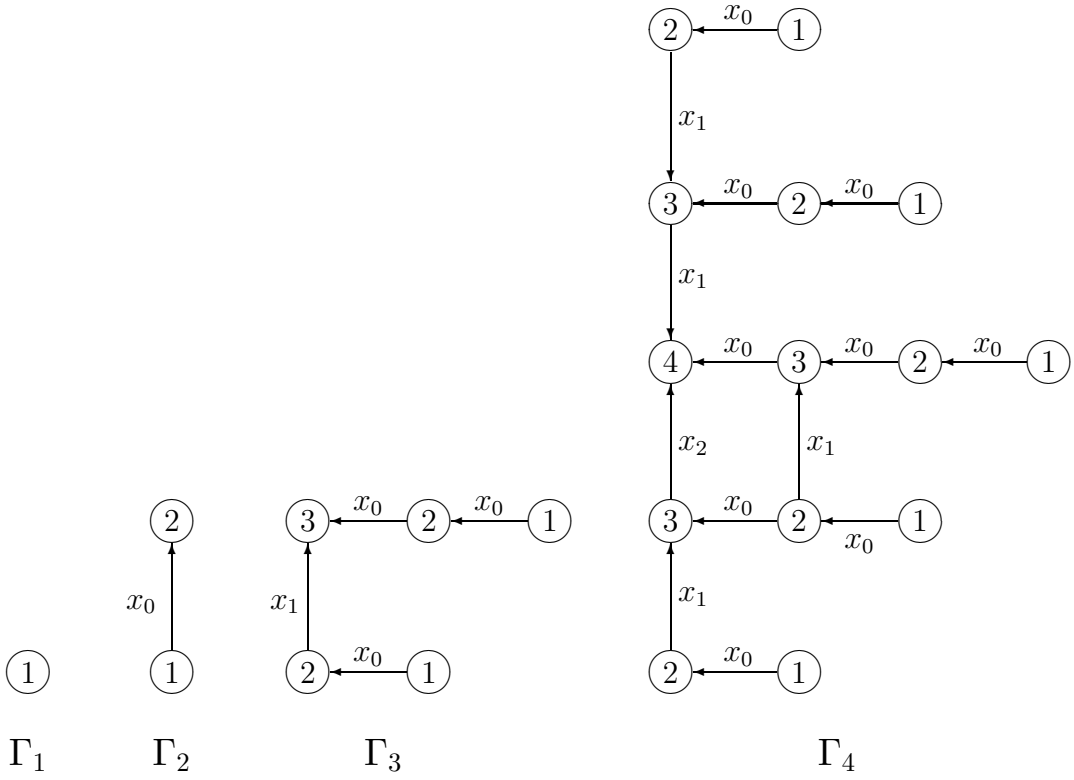

Lemma 1 *Let a_{nk} $(1 \le k \le n)$ be the number of vertices of Γ_n that have rank k. Then*

$$a_{nk} = \frac{k(2n-k-1)!}{(n-k)!\,n!}. \tag{9}$$

The total number of vertices in Γ_n equals the nth Catalan number, that is,

$$\frac{(2n)!}{n!\,(n+1)!}\,.$$

Proof. We proceed by induction on n. Obviously, $a_{11} = 1$, which agrees with the formula. Each vertex of rank k in Γ_n, where $1 \le k \le n$, becomes a vertex of rank $k+1$ in $\Psi(\Gamma_n)$. After applying $\mathcal{A}_0$ to $\Psi(\Gamma_n)$, the column of this vertex will contain exactly one vertex of each rank from 1 to $k+1$. Then the number of vertices of rank i in $\Gamma_{n+1} = \mathcal{A}_0(\Psi(\Gamma_n))$ will be equal to the sum $\sum_{i\le k+1} a_{nk}$. Therefore we have the recursive formulas $a_{n+1,1} = a_{n+1,2} = a_{n1} + \cdots + a_{nn}$, $a_{n+1,3} = a_{n2} + \cdots + a_{nn}$, $\ldots$, $a_{n+1,n} = a_{n,n-1} + a_{nn}$, $a_{n+1,n+1} = a_{nn}$. By the inductive assumption, $a_{n+1,n+1} = a_{nn} = 1$. Suppose that we have already proved formula (9) for $a_{n,k+1}$, where $2 \le k \le n$. Then

$$\begin{aligned} a_{n+1,k} &= a_{n,k-1} + a_{n+1,k+1} = \frac{(k-1)(2n-k)!}{(n-k+1)!\,n!} + \frac{(k+1)(2n-k)!}{(n-k)!\,(n+1)!} \\ &= \frac{(2n-k)!}{(n-k)!\,n!}\left(\frac{k-1}{n-k+1} + \frac{k+1}{n+1}\right) = \frac{k(2n-k+1)!}{(n-k+1)!\,(n+1)!}, \end{aligned}$$

which proves (9) for $a_{n+1,k}$. Finally,

$$a_{n+1,1} = a_{n+1,2} = \frac{2(2n-1)!}{(n-1)!\,(n+1)!} = \frac{(2n)!}{n!\,(n+1)!},$$

which proves (9) for $a_{n+1,1}$.

Note that the total number of vertices in Γ_n will be equal to $a_{n1} + \cdots + a_{nn} = a_{n+1,1} = (2n)!/(n!\,(n+1)!)$, which is the nth Catalan number. This completes the proof of the Lemma.

Now let b_{nk} will be the number of edges in Γ_k labelled by x_k, where $0 \le k \le n$. (Recall that if we denote a vertex by a circle with the number m inside, then this vertex has a loop labelled by x_m.)

Lemma 2 *For any $n \ge 2$,*

$$b_{n0} = b_{n1} = \frac{3(2n-2)!}{(n-2)!\,(n+1)!}.$$

Proof. (Remark that $b_{10} = 0$, $b_{11} = 1$.) We would like to establish a more general formula

$$b_{nk} = \frac{(k+1)(2n-k-2)!}{(n-k)!\,(n+1)!} \cdot (3n^2 - 3n(k+1) + k^2 + 2k) \tag{10}$$

for any $0 \le k < n$ together with $b_{nn} = 1$. The conclusion of the Lemma is a partial case of these. We proceed by induction on n. The case $n = 1$ is obvious. First of all, we want to give recursive formulas for $b_{n+1,i}$ $(0 \le i \le n+1)$. Let v be a vertex of rank k in Γ_n $(1 \le k \le n)$. It becomes a vertex of rank $k+1$ in $\Psi(\Gamma_n)$ and thus creates exactly k edges labelled by x_0 in $\mathcal{A}_0(\Psi(\Gamma_n)) = \Gamma_{n+1}$. This implies $b_{n+1,0} = a_{n1} + 2a_{n2} + \cdots + na_{nn}$. Using the recursive formulas from the proof of Lemma 1, one can rewrite this as $b_{n+1,0} = (a_{n1} + \cdots + a_{nn}) + (a_{n2} + \cdots + a_{nn}) + \cdots + a_{nn} = a_{n+1,2} + a_{n+1,3} + \cdots + a_{n+1,n+1} = a_{n+2,3}$. Therefore, $b_{n+1,0} = 3(2n)!/((n-3)!n!)$, which proves (10) for $b_{n+1,0}$.

Let e be an edge labelled by x_k in Γ_n $(0 \le k \le n)$. It becomes and edge labelled by x_{k+1} in $\Psi(\Gamma_n)$ so it creates one edge of each of the ranks from 1 to $k+1$ in $\mathcal{A}_0(\Psi(\Gamma_n)) = \Gamma_{n+1}$. Thus $b_{n+1,i}$ will be the sum $\sum_{i \le k+1} b_{nk}$ for any $1 \le i \le n+1$. We have the recursive formulas $b_{n+1,1} = b_{n0} + \cdots + b_{nn}$, $b_{n+1,2} = b_{n1} + \cdots + b_{nn}$, $\ldots$, $b_{n+1,n+1} = b_{nn}$. Therefore, $b_{n+1,n+1} = b_{nn} = 1$ by the inductive assumption. Suppose that (10) has been proved for $b_{n+1,k+1}$, where $1 \le k \le n$. We have $b_{n+1,k} = b_{n,k-1} + b_{n+1,k+1}$. Hence, using the assumptions, we obtain that $b_{n+1,k}$ equals

$$\frac{k(2n-k-1)!}{(n-k+1)!\,(n+1)!} \cdot (3n^2 - 3nk + k^2 - 1) + \frac{(k+2)(2n-k-1)!}{(n-k)!\,(n+2)!} \cdot (3n^2 - 3nk + k^2 + k)$$

$$= \frac{(2n-k-1)!}{(n-k)!\,(n+1)!} \left(\frac{k(3n^2 - 3nk + k^2 - 1)}{n-k+1} + \frac{(k+2)(3n^2 - 3nk + k^2 + k)}{n+2} \right)$$

$$= \frac{(k+1)(2n-k)!}{(n-k+1)!\,(n+2)!} \cdot (3n^2 - 3nk + 3n + k^2 - k),$$

which proves (10) for $b_{n+1,k}$.

The proof is complete.

Finally, to compute $\delta(\bar{\Gamma}_n)$, one has to divide the number of oriented edges in $\bar{\Gamma}_n$ labelled by $x_0^{\pm 1}$, $x_1^{\pm 1}$ by the number of vertices in Γ_n (it is the same as the number of vertices in $\bar{\Gamma}_n$). According to Lemmas 1 and 2, we have

$$\delta(\bar{\Gamma}_n) = \frac{2(b_{n0} + b_{n1})}{a_{n+1,1}} = \frac{12(2n-2)!}{(n-2)!\,(n+1)!} \cdot \frac{n!\,(n+1)!}{(2n)!} = \frac{6(n-1)}{2n-1}$$

for any $n \geq 2$.

This completes the proof. It may be interesting to know how many vertices in $\bar{\Gamma}_n$ have a given degree. For any $n \geq 5$, one has

$$\nu_2 = \frac{3(2n-4)!}{(n-2)!\,(n-1)!},\ \nu_3 = \frac{4(5n-12)(2n-5)!}{(n-3)!\,n!},\ \nu_4 = \frac{6(2n-5)!}{(n-5)!\,(n+1)!}.$$

Here ν_d denotes the number of vertices that have degree d. This implies that an average vertex of $\bar{\Gamma}_{m,n}$ has degree 2, 3, 4 with probabilities approaching 3/16, 5/8, 3/16, respectively, as m, n approach infinity.

2 The Lower Bound for the Growth Rate

By a *norm* of an element of F we mean the shortest length of a word in generators $\{x_0^{\pm 1}, x_1^{\pm 1}\}$ that represents this element. Equivalently, the norm of an element is the distance from this element to the identity in the Cayley graph $\mathcal{C}_2$ of the group F.

Let b_n ($n \geq 0$) be the number of elements in the ball around 1 in $\mathcal{C}_2$ of radius n. Since F has a free subsemigroup generated by x_0, x_1, the group F has exponential growth. But it is still unknown what will be the general growth function

$$G(t) = \sum_{n=0}^{\infty} b_n = b_0 + b_1 t + b_2 t^2 + \cdots + b_n t^n + \cdots$$

of the group F in canonical generators x_0, x_1. Since the sequence b_n is monotone and satisfy the submultiplicative inequality $b_{m+n} \leq b_m b_n$ ($m, n \geq 0$), there exists the limit

$$b = \lim_{n\to\infty} \sqrt[n]{b_n},$$

which is called the *growth rate* of a group (in a given set of generators). The exact value of the growth rate of F in generators x_0, x_1 is also unknown. Recently Burillo [5] found an exact number of **positive** elements in the ball of radius n (an element of F is positive if it is a product of positive letters x_0, x_1, x_2, ... of the infinite set of generators). He also found the general growth function for positive elements and its growth rate. This implies the lower bound for the growth rate of F. Namely, Burillo proved that b is not less than the largest root of the equation $x^3 - x^2 - 2x + 1 = 0$, that is, $b \geq 2.2469796\ldots$. We would like to present a better lower bound for the growth rate of F.

Theorem 3 *The growth rate of the group F in generators x_0, x_1 is not less than* $(3+\sqrt{5})/2 = 2.6180339\ldots$.

Proof. We use the normal form for the elements of F found in [17]. It was shown that there exists a regular spanning tree of the Cayley graph $\mathcal{C}_2$. Each element in F can be uniquely represented by a word of the form $w(x_0, x_1)$, where w does not contain the following forbidden subwords: $x_i^{\pm 1}x_i^{\mp 1}$ $(i = 0, 1)$, $x_1^{\pm 1}x_0^n x_1$, $x_1^{\pm 1}x_0^{n+1}x_1^{-1}$ $(n \geq 1)$. This normal form of an element does not give a minimal representative of an element in the sense of its norm (that is, the regular tree is not geodesic). However, for the regular language L of normal forms one can find the growth rate using standard methods. If a word of length n belongs to this language, then the norm of the corresponding element of F does not exceed n and so it will belong to $B_n(1)$. Thus b_n, the number of elements in $B_n(1)$, will be not less than the number of words in L that have length n. So we begin to calculate the growth function of the regular language L. To draw its generating automaton, we need to subdivide L into seven disjoint subsets.

1) The empty word.
2) The words of the form x_0^k, where $k \geq 1$.
3) The words that end with the letter x_0^{-1}.
4) The words that end with the letter x_1.
5) The words that end with the letter x_1^{-1}.
6) The words that end with x_1x_0 or $x_1^{-1}x_0$.
7) The words that end with $x_1x_0^2$ or $x_1^{-1}x_0^2$.

It is obvious that all the seven classes are pair-wise disjoint and their union is L. This allows to draw the generating automaton of L:

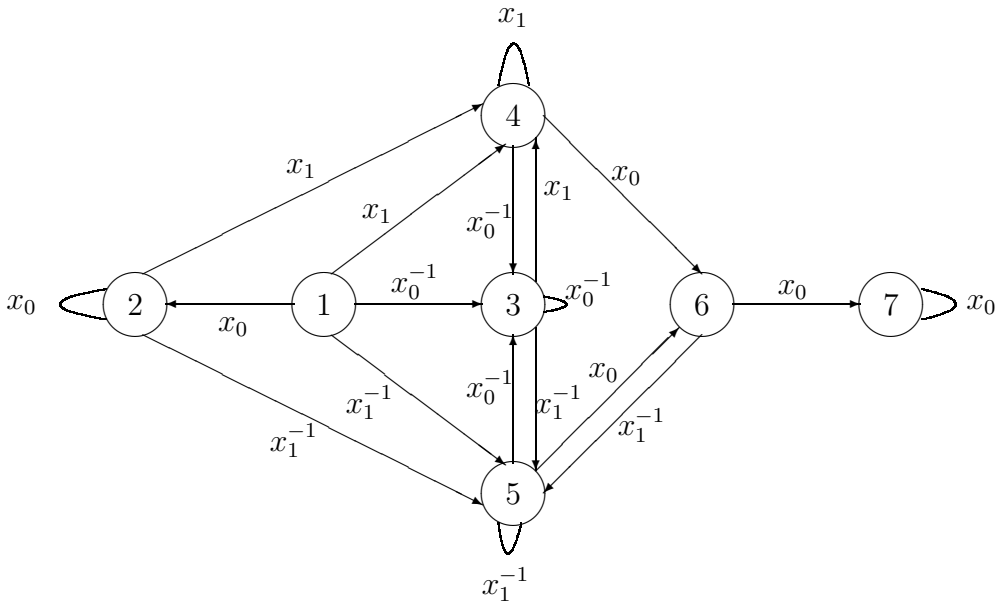

Now let A be the transition matrix of the automaton, that is, $a_{ij} = 1$ if there is a directed edge from the ith to the jth vertex of the automaton and $a_{ij} = 0$ otherwise. Obviously,

$$A = \begin{pmatrix} 0 & 1 & 1 & 1 & 1 & 0 & 0 \\ 0 & 1 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix}$$

Let $c_n^{(p)}$ be the number of directed paths of length n in the automaton from the 1st vertex to the vertex number p ($n \geq 0$, $1 \leq p \leq 7$). Clearly, $c_0^{(1)} = 1$, $c_0^{(p)} = 0$ if $2 \leq p \leq 7$. For any $n \geq 1$, $1 \leq p \leq 7$ one has

$$c_n^{(p)} = c_{n-1}^{(1)} a_{1p} + c_{n-1}^{(2)} a_{2p} + \cdots + c_{n-1}^{(7)} a_{7p}. \tag{11}$$

Let $C_p(t)$ be the generating function of $c_n^{(p)}$, that is,

$$C_p(t) = \sum_{n=0}^{\infty} c_n^{(p)} t^n = c_0^{(p)} + c_1^{(p)} t + \cdots + c_n^{(p)} t^n + \cdots$$

for each $1 \leq p \leq 7$. Obviously, from (11) we obtain $C_1 = 1$, $C_p(t) = t \sum_{i=1}^{7} C_i(t) a_{ip}$. In the matrix form, this means that $(C_1(t), \ldots, C_7(t))(E - tA) = (1, 0, \ldots, 0)$. Thus $(C_1(t), \ldots, C_7(t))$ is $(1, 0, \ldots, 0)(E - tA)^{-1}$, that is, the first row of $(E - tA)^{-1}$. We are interested in the function $C(t) = C_1(t) + \cdots + C_7(t)$. Finding the inverse matrix to $E - tA$ and adding the elements in its first row, gives us

$$C(t) = \frac{t^3 + 1}{(t-1)(t^2 - 3t + 1)}$$

which is presented by series as

$$C(t) = 1 + 4t + 12t^2 + 34t^3 + 92t^4 + 244t^5 + 642t^6 + 1684t^7 + 4412t^8 + 11554t^9 + \cdots. \tag{12}$$

The coefficient c_n on t^n shows the number of directed paths of length n in the automaton that start in the first vertex. Hence $C(t)$ is the generating function for the regular language L. The radius of convergence of the series (12) will be the least absolute value of a root of the denominator, that is, it will be $(3 - \sqrt{5})/2$. So the growth rate of L will be the reciprocal, that is, $(3 + \sqrt{5})/2$. We have already mentioned that $\sum_{i=0}^{n} c_n \leq b_n$ for all $n \geq 0$. So we have the lower bound for the growth function of F and $(3 + \sqrt{5})/2 = 2.6180339\ldots$ will be the lower bound for its growth rate. This completes the proof.

To compare the growth function of L with the growth function of F, let us show the first terms of the series for the function $G_s(t)$, where the coefficient on t^n below shows the number $s_n = b_n - b_{n-1}$, that is, the number of elements in the sphere of radius n in the Cayley graph of F. It is easy to see that the growth rates of b_n and s_n are the same. The coefficients of the series below were found on a computer.

$$G_s(t) = 1 + 4t + 12t^2 + 36t^3 + 108t^4 + 314t^5 + 906t^6 + 2576t^7 + 7280t^8 + 20352t^9 + \cdots.$$

One can mention that the sequence s_n/s_{n-1} is decreasing for $1 \leq n \leq 9$. Since $s_9/s_8 = 2.7956043\ldots$, one can expect that the growth rate of F has this number as an upper bound. (Of course, this is only a conjecture.)

3 Representing Elements of F by Diagrams

This background Section is devoted to the description of the representation of elements of F by semigroup diagrams. We need this for the result of the next Section. The contents of the present Section is essentially known. Detailed information about diagram groups can be found in [18]. However, we need to describe a modified version of this idea based on the representation of F by non-spherical diagrams.

First of all, let us recall the concept of a semigroup diagram and introduce some notation. To do this, we consider the following example. Let $\mathcal{P} = \langle\, a, b \mid aba = b, bab = a \,\rangle$ be the semigroup presentation. It is easy to see by the following algebraic calculation

$$a^5 = a(bab)a(bab)a = (aba)(bab)(aba) = bab = a$$

that the words a^5 and a are equal modulo $\mathcal{P}$. The same can be seen from the following picture

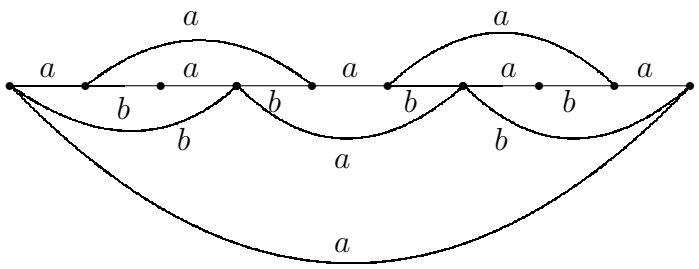

This is a *diagram* Δ over the semigroup presentation $\mathcal{P}$. It is a plane graph with 10 vertices, 15 (geometric) edges and 6 faces or *cells*. Each cell corresponds to an elementary transformation of a word, that is, a transformation of the form $p \cdot u \cdot q \to p \cdot v \cdot q$, where p, q are words (possibly, empty), $u = v$ or $v = u$ belongs to the set of defining relations. The diagram Δ has the leftmost vertex denoted by $\iota(\Delta)$ and the rightmost vertex denoted by $\tau(\Delta)$. It also has the *top path* $\mathbf{top}(\Delta)$ and the *bottom path* $\mathbf{bot}(\Delta)$ from $\iota(\Delta)$ to $\tau(\Delta)$. Each cell π of a diagram can be regarded as a diagram itself. The above functions ι, τ, $\mathbf{top}$, $\mathbf{bot}$ can be applied to π as well. We do not distinguish isotopic diagrams.

We say that Δ is a (w_1, w_2)-diagram whenever the label of its top path is w_1 and the label of its bottom path is w_2. In our example, we deal with an (a^5, a)-diagram. If we have two diagrams such that the bottom path of the first of them has the same label as the top path of the second, then we can naturally *concatenate* these diagrams by identifying the bottom path of the first diagram with the top path of the second diagram. The result of the concatenation of a (w_1, w_2)-diagram and a (w_2, w_3)-diagram obviously is a (w_1, w_3)-diagram. We use the sign $\circ$ for the operation of concatenation. For any diagram Δ over $\mathcal{P}$ one can consider its *mirror image* Δ^{-1}. A diagram may have *dipoles*, that is, subdiagrams of the form $\pi \circ \pi^{-1}$, where π is a single cell. To *cancel* (or *reduce*) the dipole means to remove the common boundary of π and π^{-1} and then to identify $\mathbf{top}(\pi)$ with $\mathbf{bot}(\pi^{-1})$. In any digaram, we can cancel all its dipoles, step by step. The result does not depend on the order of cancellations. A diagram is *irreducible* whenever it has no dipoles. The operation of cancelling dipoles has an inverse operation called the *insertion* of a dipole. These operations

induce an equivalence relation on the set of diagrams (two diagrams are *equivalent* whenever one can go from one of them to the other by a finite sequence of cancelling/inserting dipoles). Each equivalence class contains exactly one irreducible diagram.

For any nonempty word w, the set of all (w, w)-diagrams forms a monoid with the identity element $\varepsilon(w)$ (the diagram with no cells). The operation $\circ$ naturally induces some operation on the set of equivalence classes of diagrams. This operation is called a *product* and equivalent diagram are called *equal*. (The sign $\equiv$ will be used to denote that two diagrams are isotopic.) So the set of all equivalence classes of (w, w)-diagram forms a group that is called the *diagram group* over $\mathcal{P}$ with *base* w. We denote this group by $\mathcal{D}(\mathcal{P}, w)$. We can think about this group as of the set of all irreducible (w, w)-diagrams. The group operation is the concatenation with cancelling all dipoles in the result. An inverse element of a diagram is its mirror image. We also need one more natural operation on the set of diagrams. By the *sum* of two diagrams we mean the diagram obtained by identifying the rightmost vertex of the first summand with the leftmost vertex of the second summand. This operation is also associative. The sum of diagrams Δ_1, Δ_2 is denoted by $\Delta_1 + \Delta_2$.

It is known that the group F is the diagram group over the simplest semigroup presentation $\mathcal{P} = \langle\, x \mid x^2 = x \,\rangle$ with base x (note that for any base x^k, where $k \geq 1$, we get the same group). It is not hard to compare this representation of F with other known representations. Say, in [6] and many other papers, the elements of F are represented by certain pairs of rooted binary trees. If an element $g \in F$ is represented by an (irreducible) (x, x)-diagram Δ, then the corresponding pair of trees can be recovered as follows. Let us consider the dual graph Γ of the graph Δ. The vertices of Γ are the midpoints of all edges of Δ. For each (x, x^2)-cell of Δ we connect the midpoint of the top edge of the cell (labelled by x) with the midpoints of the edges e', e'', where $e'e''$ is the bottom path of the cell. This gives a "wedge" usually called the *caret*. All these carets form a rooted binary tree. Similarly, we do the same thing for all (x^2, x)-cells of Δ. This gives the other rooted binary tree. Each of the two trees have the same number of *leaves*, that is, the vertices which are not roots of any carets. The ith leaf of the upper tree (if to count the leaves from the left to the right) coincides with the ith leaf of the lower tree.

Now let us compare the diagram representation of F with the representation of its elements by piecewise-linear homeomorphisms of the closed unit interval $[0, 1]$. Let Δ be an (x^p, x^q)-diagram over $\mathcal{P}$. We will show how to assign to it a piecewise-linear function from $[0, p]$ onto $[0, q]$. Each positive edge of Δ is homeomorphic to the unit interval $[0, 1]$. So we assign a coordinate to each point of this edge (the leftmost end of an edge has coordinate 0, the rightmost one has coordinate 1). Let π be an (x, x^2)-cell of Δ. Let us map $\mathbf{top}(\pi)$ onto $\mathbf{bot}(\pi)$ linearly, that it, the point on the edge $\mathbf{top}(\pi)$ with coordinate $t \in [0, 1]$ is taken to the point on $\mathbf{bot}(\pi)$ with coordinate $2t$ (the bottom path of π has length 2 so it is naturally homeomorphic to $[0, 2]$). The same thing can be done for an (x^2, x)-cell of Δ. Thus for any cell π of Δ we have a natural mapping T_π from $\mathbf{top}(\pi)$ onto $\mathbf{bot}(\pi)$ (we call it a *transition map*). Now let t be any number in $[0, p]$. We consider the point o on $\mathbf{top}(\Delta)$ that has coordinate t. If o is not a point of $\mathbf{bot}(\Delta)$, then it is an internal point on the top path of some cell. Thus we can apply the corresponding transition map to o. We repeat this operation until we get a point o' on the path $\mathbf{bot}(\Delta)$. The coordinate of this point is a number in $[0, q]$. Hence we have a function $f_\Delta \colon [0, p] \to [0, q]$ induced by Δ. It is easy to see

this will be a piecewise-linear function. When we concatenate diagrams, this corresponds to the composition of the PL functions induced by these diagrams. For the group F, which is the diagram group $\mathcal{D}(\mathcal{P}, x)$, we have the homeomorphism from it to $PLF[0,1]$. It is known this is an isomorphism.

Now let us recall that F has another representation by PL functions. Namely, let $f_i\colon [0,\infty) \to [0,\infty)$ ($i \geq 0$ is an integer) be the PL function that has slope 1 on $[0,i]$, slope 2 on $[i,i+1]$ and slope 1 again on $[i+1,\infty)$. It is easy to verify that for any integers $j > i \geq 0$, one has $f_j f_i = f_i f_{j+1}$ (the functions act on the right). Thus the mapping $x_i \mapsto f_i$ induces a homomorphism from F to $PLF[0,\infty)$. It is also known this is an isomorphism.

In the above representation of F by diagrams, all these diagrams were *spherical*, that is, they were (w,w)-diagrams for some word w. Our new representation of F by non-spherical diagrams, which we are going to describe, corresponds to the representation of F by the PL functions on $[0,\infty)$. Let Δ be any diagram over $\mathcal{P} = \langle\, x \mid x^2 = x \,\rangle$, not necessarily spherical. Let us add an infinite sequence of edges on the right of Δ, each edge is labelled by x. This object will be called an *infinite diagram* over $\mathcal{P}$. Note that it has finitely many cells. An infinite diagram that corresponds to Δ will be denoted by $\hat{\Delta}$. It has the leftmost vertex $\iota(\hat{\Delta})$ and two distinguished infinite paths starting at $\iota(\hat{\Delta})$, both labelled by the infinite power of x. These paths will be denoted by $\mathbf{top}(\hat{\Delta})$ and $\mathbf{bot}(\hat{\Delta})$, respectively. The concept of a dipole in an infinite diagram is defined as above. The same concerns the operations of deleting/inserting a dipole, the equivalence relation induced by these operations, and so on. Any two infinite diagrams can be naturally concatenated (the bottom path of the first factor is identified with the top path of the second factor). We use the sign $\circ$ for the concatenation. The operation $\circ$ gives the set of all infinite diagrams a monoid structure. The identity of it is the infinite diagram without cells denoted by ε. As in the case of spherical diagrams, the operation of concatenation induces a group operation on the set of all equivalence classes of infinite diagrams. Thus we have a group. We shall denote it by $\hat{\mathcal{D}}(\mathcal{P}, x)$. (This makes sense for any semigroup presentation $\mathcal{P}$ in an alphabet of one letter. Note that we can forget about the labels if we work with a one-letter alphabet.) It is quite easy to see that the group we have will be isomorphic to F. Indeed, let X_i be the infinite diagram

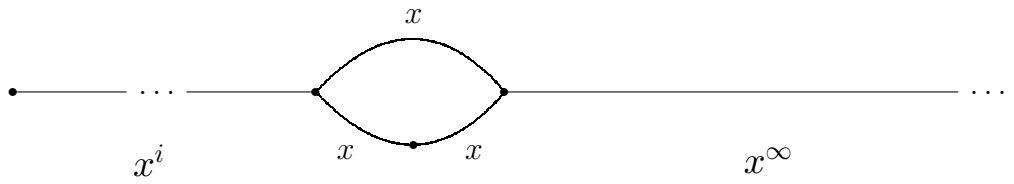

that consists of an (x, x^2)-cell, the finite path labelled by x^i on the left of it and the infinite path labelled by the infinite power of x, on the right. By X_i^{-1} we mean the mirror image of X_i under the horizontal axis symmetry. Infinite diagrams of the form $X_i^{\pm 1}$ ($i \geq 0$) are called *atomic*. For any integers $j > i \geq 0$, the diagram

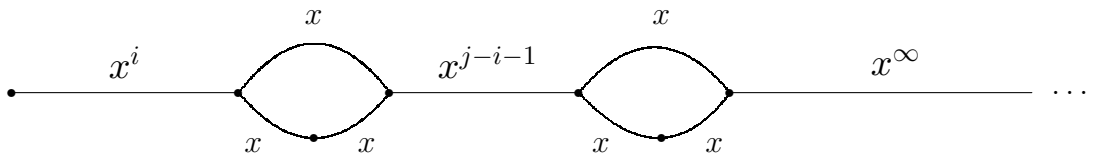

equals both $X_j \circ X_i$ and $X_i \circ X_{j+1}$. This means that we have a homomorphism from F to the group of infinite diagrams. This homomorphism is onto because any infinite diagram is a concatenation of atomic diagrams. The homomorphism must be injective because all homomorphic images of F are abelian [6]. However, the group of infinite diagrams is not abelian since $X_1X_0 = X_0X_2 \neq X_0X_1$.

Note that the diagram X_i corresponds to the above defined PL function f_i. (The transition maps allow to assign a PL function to any infinite diagram.) If we consider a dual graph to an infinite diagram, this gives us a pair of binary rooted forests. Each forest is a sequence of binary rooted trees. All of these trees but the finite number of them consists of just the root vertex. All leaves of a forest can be naturally enumerated. Given a pair of such forests, we can identify the ith leaf of the "upper" forest with the ith leaf of the "lower" forest for all $i \geq 0$.

When we work with infinite diagrams, it is convenient to eliminate the infinite "tail" on the right of each infinite diagram. An ordinary diagram over $\mathcal{P}$ is called *canonical* whenever it has no dipoles and it is not a sum of a diagram and an edge. It is obvious that there is a one-to-one correspondence between the set of infinite diagrams without dipoles and the set of canonical diagrams. So we may assume that each element of F has a unique canonical representative. We have a group structure on the set of all canonical diagrams over $\mathcal{P}$. Given an (x^p, x^q)-diagram Δ_1 and an (x^s, x^t)-diagram Δ_2, we multiply them as follows. If $q = s$, then we concatenate them. If $q < s$, then we concatenate $\Delta_1 + \varepsilon(x^{s-q})$ and Δ_2. If $q > s$, then we concatenate Δ_1 and $\Delta_2 + \varepsilon(x^{q-s})$. After the concatenation, we reduce all dipoles in the result. Then we need to make the diagram canonical. This means that we have to delete the rightmost common suffix of the top and the bottom path of the diagram we have. The only exception is made for the diagram $\varepsilon(x)$, the identity element of the group. This diagram is already canonical so we leave it as it is.

This representation of F by canonical diagrams is very convenient for many reasons. Say, the number of cells in the canonical diagram of an element always equals the length of the normal form of this element. Given a normal form, it is very easy to draw the corresponding diagram, and vice versa. The following example illustrates the diagram that corresponds to the element $g = x_0^2 x_1 x_6 x_3^{-1} x_0^{-2}$ represented by its normal form:

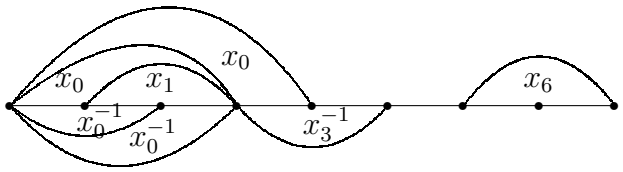

This representation of elements in F by canonical diagrams has some more advantages. First of all, the number of cells in a canonical diagram of an element g is obviously the length of the normal form of g, which is also the shortest length of a word in the infinite set of generators $\{ x_0^{\pm 1}, x_1^{\pm 1}, x_2^{\pm 1}, \ldots \}$ that represents g. Given the canonical representative Δ of G, it is easy to construct an (x, x)-diagram that represents the same element. One has to add an edge on the right of Δ, that is, to consider the diagram $\Delta + \varepsilon(x)$. Let v be the rightmost vertex of this diagram. We then connect v by arcs to all the vertices in $\mathbf{top}(\Delta)$

in the upper part and connect v by arcs to all the vertices of Δ in the lower part. All edges we add have label x. As a result, we give an (x, x)-diagram that represents g in the diagram group $\mathcal{D}(\mathcal{P}, x)$. This diagram has "redundant" cells — they are all cells that are not in Δ. These cells do not correspond to any generators so Δ looks a more "natural" representative of g.

If we prefer to work with dual graphs, then, given a pair of binary rooted trees from [6] that represent an element g, one has to remove all right carets from it. The rest will be the pair of forests that is exactly the dual graph to the canonical representative of g.

Note that the diagram in the above picture is drawn in such a way that all the (x, x^2)-cells of it are above the straight line and all the (x^2, x)-cells are below the straight line. This is an easy consequence of the fact that the normal form of an element in F is a product of the form pq^{-1}, where p, q are positive words (words in generators x_0, x_1, x_2, ... with no negative exponents). The following elementary fact was essentially used several times in [19, 16] and some other papers.

Lemma 3 *Let $\mathcal{P} = \langle X \mid \mathcal{R} \rangle$ be a semigroup presentation. Suppose that all defining relations of $\mathcal{P}$ have the form $a = A$, where $a \in X$, A is a word of length at least 2. Also assume that all letters in the left-hand sides of the defining relations are different. Then any irreducible diagram Δ over $\mathcal{P}$ is the concatenation of the form $\Delta_1 \circ \Delta_2^{-1}$, where the top path of each cell of both Δ_1, Δ_2 has length 1. The longest positive path in Δ from $\iota(\Delta)$ to $\tau(\Delta)$ is the bottom path of both Δ_1 and Δ_2.*

Note that $\langle x \mid x = x^2 \rangle$ obviously satisfy the conditions of the Lemma. The same concerns the presentation $\langle a, b \mid a = bab, b = aba \rangle$, which was considered in the previous Section. Let us recall the idea of the proof. Let p be the longest positive path in Δ from $\iota(\Delta)$ to $\tau(\Delta)$. It cuts Δ into two parts. It suffices to prove that all cells in the "upper" part correspond to the defining relations of the form $a = A$, where a is a letter, and none of them correspond to $A = a$. Assume the contrary. Suppose that there is a cell π in the upper part of Δ with the top label A and the bottom label a. The bottom path of π cannot be a subpath in p since p is chosen the longest. So the bottom edge of π belongs to the top path of some cell π'. The diagram Δ has no dipoles. All letters in the left-hand sides of the defining relations are different. So the top path of π' cannot have length 1. This means that we have found a new cell in the upper part of Δ that also corresponds to the defining relation of the form $A = a$. Applying the same argument to π', we get a process that never terminates. This is impossible since the cells that appear during the process cannot repeat. This completes the proof.

4 The Length Formula

An algorithm to find the norm of an element of F in generators x_0, x_1 was obtained by Fordham [12] in his PhD. The algorithm is based on the representation of F by pairs of binary rooted trees. Each caret of a tree belongs to one of 7 classes defined in [12] (only 6 of these classes are essential). All carets in each of the trees are enumerated. For any number, there are 36 cases for a pair formed by the ith carets of the trees, to belong to some of the

classes. To each of these 36 cases one can assign a weight. The sum of these numbers will be the norm of an element. By this algorithm, the norm can be found in a polynomial time, although the algorithm itself is quite complicated. However, it is a powerful tool to solve various problems about F. For instance, Fordham's algorithm is successfully implemented in recent papers [5, 8]. In this Section, we present another algorithm to find the norm of an element in F. We are based on the representation of F by non-spherical diagrams described in Section 3. We classify vertices of the canonical diagram of a given element. There are only two cases for this. To find the norm, one has to add the number of cells in the diagram and the number of so-called special vertices multiplied by 2. This algorithm is very clear and easy in use.

Let Δ be a diagram over $\mathcal{P} = \langle\, x \mid x^2 = x \,\rangle$. From Lemma 3 we know that all vertices of Δ belong to the longest positive path p from $\iota(\Delta)$ to $\tau(\Delta)$. We need to classify vertices of Δ. A non-empty positive path q in Δ (in particular, a single edge) is called a *brigde* whenever q is a common subpath in both $\mathbf{top}(\Delta)$ and $\mathbf{bot}(\Delta)$. If q is a bridge of Δ, then one has a decomposition of the form $\Delta' + q + \Delta''$, where Δ, Δ'' are some diagrams (possibly Δ'' is empty) and q denotes the subdiagram in Δ that consists of the bridge q. We say that the bridge q is *nontrivial* whenever Δ'' has cells. Note that any bridge q, being a path, has its initial point $\iota(q)$.

A vertex v in Δ is called *active* if v is either an initial point of a cell in Δ or v is an initial point of a nontrivial bridge of Δ. Since Δ is a finite graph, we have the distance function on vertices of Δ. The distance between two vertices is the length of the shortest path in Δ that connects these vertices. An active vertex in Δ is called *special* if its distance from the origin $O = \iota(\Delta)$ is strictly greater than 1.

In the following picture

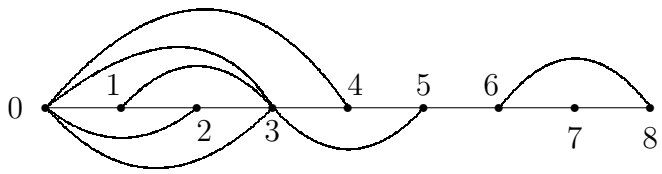

we enumerate all vertices of the diagram Δ travelling along p from the origin and starting from the number 0. It is clear that vertices 0, 1, 3, 5, 6 are active. Vertices 5, 6 from this list are the only special vertices of Δ.

Theorem 4 *Let $g \in F$ be represented by the canonical diagram Δ. Then the norm of g, that is, the length of the shortest word representing g in group generators x_0, x_1, can be found by the following formula:*

$$||g|| = \#_c\Delta + 2\#_s\Delta, \tag{13}$$

where $\#_c$ denotes the number of cells and $\#_s$ denotes the number of special vertices of a diagram.

In the above example, one easily has $\#_c\Delta = 7$, $\#_s\Delta = 2$ so $||g|| = ||x_0^2x_1x_6x_3^{-1}x_0^{-2}|| = 7 + 2 \cdot 2 = 11$. Recall that the number of cells in the canonical diagram of g is always equal to the length of the normal form of g. Before the proof, we would like to verify that g is an example of a *dead vertex* in the Cayley graph of F. Namely, we would like to show that for any $a = x_i^{\pm 1}$ $(i = 0, 1)$ one has $||ga|| < ||g||$ (that is, the norm of g decreases in all directions). It can be shown that the number 11 is the smallest one for which this effect is possible.

We need to analyze what happens if we multiply g by each of the four semigroup generators. Let us consider 4 cases.

1) $g \to gx_0$. In this transition, when we multiply g by x_0, one cell in Δ cancels (the leftmost cell on the bottom). The set of vertices remains the same. It is also clear that the set of active vertices will be the same and no vertices change their distance from the origin. So $\#_c$ decreases by 1, $\#_s$ is the same. Hence the norm decreases by 1, that is, $||gx_0|| = ||g|| - 1 = 10$.

2) $g \to gx_1$. When we multiply g on the right by x_1, we cancel the second cell on the bottom. This means we just remove the arc that connect vertices 3 and 5. In this case the vertex number 3 will be no longer active. But the vertex number 4 becomes active because it will be the initial point of a bridge. However, the vertex number 4 is on the distance 1 from the origin so it is not special. Thus the only special vertices in the new diagram will be 5 and 6, as before. This means that the norm decreases by 1 so $||gx_1|| = ||g|| - 1 = 10$.

3) $g \to gx_0^{-1}$. Multiplying by x_0^{-1} on the right means that we add a new cell to the diagram. This means we add an arc that connects vertices 0 and 5. The set of active vertices remains the same. But now the vertex number 5 will be on the distance 1 from the origin so the only special vertex in the new diagram will be the vertex number 6. Thus $\#_c$ increases by 1, $\#_s$ decreases by 1 and so the norm decreases by 1, that is, $||gx_0^{-1}|| = ||g|| - 1 = 10$.

4) $g \to gx_1^{-1}$. In this case we need to add a new cell connecting vertices 3 and 6 by an arc. Now the vertex number 5 will not be an initial point of a bridge. The only special vertex in the new diagram will be 6. As in the previous paragraph, $\#_c$ increases by 1, $\#_s$ decreases by 1 and $||gx_1^{-1}|| = ||g|| - 1 = 10$.

Proof of Theorem 4. We use the idea of [12, Lemma 2.1.1]. For any $g \in F$ we consider the canonical non-spherical diagram Δ that represents g. The value $\#_c\Delta + 2\#_s\Delta$ will be denoted by $\varphi(g)$. Our aim is to show that $||g|| = \varphi(g)$ for any $g \in F$.

Clearly, $\varphi(g) = 0$ if and only if $g = 1$ in F. We show that $\varphi(ga) = \varphi(g) \pm 1$ for any $g \in F$ and for any semigroup generator $a \in \{x_0^{\pm 1}, x_1^{\pm 1}\}$. This will imply that $\varphi(g) \leq ||g||$ for any $g \in F$. Indeed, if we decompose g into a product of $n = ||g||$ generators, then one can easily show by induction that $\varphi(g) \leq n$.

Let us consider two neighbour elements in the Cayley graph $\mathcal{C}_2$. For these elements we can compare the number of cells in canonical diagrams that represent these elements. It is obvious that the difference between the numbers of cells will be always 1 or -1 so we can assume without loss of generality that the number of cells in the canonical diagram of ga will be greater than the number of cells in the canonical diagram Δ of g by 1.

We consider 4 cases. In each of them we need to show that the number of special vertices either remains the same or it decreases by 1. This will imply that the value of φ always increases by 1 or decreases by 1 after we multiply its argument by a generator. Recall that

X_i is the diagram that corresponds to the generator $x_i \in F$ ($i \geq 0$).

Case 1. $a = x_0$. Since multiplying by X_0 increases the number of cells, we know that Δ has no cells whose bottom path coincides with the leftmost edge of $\mathbf{bot}(\Delta)$. When we concatenate Δ and X_0, then a new vertex appears on the bottom path of the cell of X_0. This vertex is on the distance 1 from the origin. So the vertex we add is not special. It is easy to see that no new special vertices can appear. So φ increases by 1 in this case.

Case 2. $a = x_1$. In this case we concatenate Δ and X_1. This gives a new vertex (on the bottom of the cell of X_1). Clearly, this vertex is not special since it is not an initial point of a cell and it is not the initial point of a bridge. The status of all other vertices will be the same, that is, no special vertices appear or disappear. Thus φ increases by 1.

Case 3. $a = x_0^{-1}$. Let $v_0, v_1, v_2, \ldots$ be all vertices of the bottom path of Δ enumerated from the left to the right. Multiplying by X_0^{-1} means that we connect vertices v_0 and v_2 by an arc. Suppose that v_2 was a special vertex in Δ. In this case the number of special vertices in $\Delta \circ X_0^{-1}$ will be less by 1 than the number of special vertices in Δ. Otherwise the number of special vertices remains the same. Hence φ decreases or increases by 1, respectively.

Case 4. $a = x_1^{-1}$. In the notation of the previous paragraph, one has to add an arc that connects vertices v_1 and v_3 in order to concatenate Δ and X_1^{-1}. Note that vertices v_0, v_1 are not special in our diagrams. The vertex v_3 also does not change its status because the new arc does not connect it with the origin. Clearly, nothing happens with vertices v_k if $k \geq 4$. So the only vertex we need to think about is v_2. If the distance from it to the origin in Δ is 1, then the same is true for $\Delta' = \Delta \circ X_1^{-1}$. Suppose that the distance in Δ from v_2 to the origin is at least 2. Then the same is true for Δ'. If v_2 was not active in Δ, then it could not be active in Δ'. Suppose that v_2 is an initial point of a cell in Δ. In this case v_2 be a special vertex of Δ. It is obvious that v_2 has the same status in Δ'. So the only case we need to check is when v_2 is an initial point of a nontrivial bridge in Δ. After we concatenate Δ and X_1^{-1}, this will no longer be true. So v_2 will not be special in Δ' so the number of special vertices decreases by 1. As a result, we have that φ increases or decreases by 1.

To complete the proof, we need to check that for any $g \in F$, $g \neq 1$, there exists at least one generator $a \in \{x_0^{\pm 1}, x_1^{\pm 1}\}$ such that φ decreases at its direction, that is, $\varphi(ga) = \varphi(g) - 1$. This will imply the inequality $||g|| \leq \varphi(g)$ for any $g \in F$. Indeed, if $\varphi(g) = n$, then we can start from the vertex g in $\mathcal{C}_2$. At each step we choose a direction that decreases φ by 1. After n steps, we get to the element with the zero value of φ, that is, to the identity. So the distance from g to the origin (that is, the norm of g) does not exceed n.

We enumerate the vertices of $\mathbf{bot}(\Delta)$ as above, where Δ is the canonical representative of g. Since $g \neq 1$, the diagram Δ is nontrivial. Suppose that v_2 is a special vertex in Δ. Then Δ cannot be the concatenation of some diagram Δ' and X_0 (otherwise v_2 is on the distance 1 from the origin). So if we multiply Δ by X_0^{-1}, then the number of cells increases by 1 and we are in the situation of Case 3. Since v_2 is special vertex of Δ, the value of φ decreases by 1, that is, $\varphi(gx_0^{-1}) = \varphi(g) - 1$.

Assume that v_2 is not a special vertex of Δ. Let Δ be a concatenation of some diagram Δ' and X_1^{-1}. If we multiply g by x_1 on the right, then the number of cells decreases and we are in the situation, which is inverse to the one described in Case 4 (now Δ is obtained from Δ' by adding a cell when multiplying by x_1^{-1}). We know that $\varphi(\Delta')$ is always greater by 1 than $\varphi(\Delta)$ except for one case. Let us describe it. By π we denote the cell in Δ that

corresponds to the factor X_1^{-1} in the concatenation $\Delta' \circ X_1^{-1}$. The top path of π is a product of two edges $e'e''$. The terminal point of e' (=the initial point of e'') will be denoted by v. According to Case 4 (note that we have to exchange the rôles of Δ and Δ') we have that v must be the initial point of a nontrivial bridge in Δ'. This bridge starts with the edge e''. So the terminal point of e'' must be a special vertex of Δ' and also of Δ. But this vertex is v_2, which gives a contradiction.

Let us analyze all cases that can happen to the edge e that is contained in $\mathbf{bot}(\Delta)$ (or its continuation) and has the endpoints v_1, v_2. Suppose that e does not belong to a cell so it is a bridge. This bridge should be trivial since v_2 is not special. Thus Δ is in fact an (x^m, x)-diagram for some $m \geq 1$. It is nontrivial so Δ is a concatenation of some diagram Δ' and X_0^{-1}. If we look at the situation of Case 3 (in our case we go from Δ' to Δ adding a cell), then we see that $\varphi(\Delta)$ will be always greater than $\varphi(\Delta')$ by 1 except for one case when the vertex v_1 is a special vertex of the subdiagram Δ'. This cannot happen because $\mathbf{bot}(\Delta)$ has length 1. Thus $\varphi(gx_0) = \varphi(g) - 1$ in this case.

Now let e be contained in some cell of Δ. The only case we need to consider is the one when e is a part of the bottom path of an (x, x^2)-cell (the case of an (x^2, x)-cell means that Δ is a concatenation of some diagram and X_1^{-1}, which has been already considered). Let us denote this (x, x^2)-cell by π. The edge e may be the first or the second edge of $\mathbf{bot}(\pi)$ so we have two subcases.

Suppose that $\mathbf{bot}(\pi) = e'e$ for some edge e'. If e' is contained in the bottom path of Δ, then it coincides with the first edge of $\mathbf{bot}(\Delta)$. Thus Δ will be the concatenation of some diagram Δ' and X_0. From Case 1 we can extract that $\varphi(gx_0^{-1}) = \varphi(g) - 1$ since Δ is obtained from Δ' by adding one cell that corresponds to X_0. If e' is not contained in $\mathbf{bot}(\Delta)$, then the first edge of $\mathbf{bot}(\Delta)$ must be a bottom edge of an (x^2, x)-cell. This means that Δ is a concatenation of some diagram Δ' and X_0^{-1}. The situation we have is described in Case 3. Adding the new cell to Δ' that corresponds to X_0^{-1} always increases φ except for the case when the vertex v_1 will be a special vertex of the subdiagram Δ'. But we know that the only positive edge of Δ that comes out of v_1 is e. So v_1 is not an initial point of a cell. Neither e is a bridge. So $\varphi(gx_0) = \varphi(g) - 1$ in this subcase.

Finally, let $\mathbf{bot}(\pi) = ee''$ for some edge e''. If e'' is also contained in $\mathbf{bot}(\Delta)$, then Δ is a concatenation of some diagram and X_1. From Case 2 we can easily conclude that $\varphi(gx_1^{-1}) = \varphi(g) - 1$. If e'' is not contained in $\mathbf{bot}(\Delta')$, then there is an "angle" between e'' and the third edge of $\mathbf{bot}(\Delta)$. Then v_2 will be an initial point of some cell. However, v_2 cannot be on the distance 1 from the origin because otherwise the arc that connects the origin and v_2 would cross the edge $\mathbf{top}(\pi)$, which is impossible. Therefore, v_2 is special. We have a contradiction.

The proof is complete. Note that we have a quick procedure to find minimal representatives of elements of F using the above description.

References

[1] S. I. Adian. Random walks on free periodic groups. Math. USSR Izvestiya 21:3 (1983), 425–434 (Russian original: Izv. Akad. Nauk SSSR Ser. Mat. 46 (1982) no. 6, 1139–1149,1343.).

[2] M. G. Brin, C. C. Squier. Groups of piecewise linear homeomorphisms of the real line. Invent. Math. 79 (1985), 485–498.

[3] K. S. Brown. Finiteness properties of groups. J. of Pure and Appl. Algebra, 44 (1987), 45–75.

[4] K. S. Brown, R. Geoghegan. An infinite-dimentional torsion-free FP_∞ group. Invent. Math. 77 (1984), 367–381.

[5] J. Burillo. Growth of positive words in Thompson's group F (preprint).

[6] J. W. Cannon, W. J. Floyd and W. R. Parry. Introductorary notes on Richard Thompson's groups. L'Enseignement Mathématique (2) 42 (1996), 215–256.

[7] T. Ceccherini-Silberstein, R. Grigorchuk, and P. de la Harpe. Amenability and paradoxal decompositions for pseudogroups and for discrete metric spaces. Proc. Steklov Inst. Math. 1999, no. 1 (224), 57–97 (Russian original: Trudy Mat. Inst. Steklova 224 (1999), Algebra, Topol., Differ. Uravn. i ikh Prilozh., 68–111).

[8] S. Cleary, J. Taback. Thompson's group F is not almost convex (preprint).

[9] M. Day. Amenable semigroups. Illinois J. Math. 1 (1957), 509–544.

[10] W. A. Deuber, M. Simonovitz, and V. T. Sós. A note on paradoxal metric spaces. Studia Math. 30 (1995), 17–23.

[11] E. Følner. On groups with full Banach mean value. Math. Scand. 3 (1955), 243–254.

[12] S. Blake Fordham. Minimal Length Elements of Thompson's Group F. PhD Thesis, Brigham Young University, 1995.

[13] F. P. Greenleaf. Invariant means on topological groups and their applications. Van Nostrand, 1969.

[14] R. I. Grigorchuk. On Burnside's problem on periodic groups. Funktsional. Anal. i Prilozhen. 14 (1980), 53–54 (in Russian).

[15] R. I. Grigorchuk. An example of a finitely presented amenable group which does not belong to the class EG. Sb. Math. 189 (1998) no. 1–2, 75–95 (Russian original: Mat. Sb. 189 (1998), no. 1, 79–100).

[16] V. S. Guba. Polynomial Isoperimetric Inequalities for Richard Thompson's Groups F, T, and V. In: *Algorithmic Problems in Groups and Semigroups*, J.-C. Birget et. al., Eds. Birkhäuser, Boston–Basel–Berlin, 2000, 91–120.

[17] V. S. Guba, M. V. Sapir. The Dehn function and a regular set of normal forms for R. Thompson's group F. J. Austral. Math. Soc. (Ser. A) 62 (1997), 315–328.

[18] V. S. Guba, M. V. Sapir. Diagram groups. Memoirs of the Amer. Math. Soc. 130, N 620, 1997, 1–117.

[19] V. S. Guba, M. V. Sapir. On subgroups of R. Thompson's group F and other diagram groups. Matem. Sb. 190: 8 (1999), 3–16 (Russian). English transl. in: Sbornik: Mathematics 190: 8 (1999), 1077–1130.

[20] H. Kesten. Symmetric random walks on groups. Trans. Amer. Math. Soc. 92 (1959), 336–354.

[21] J. von Neumann. Zur allgemeinen Theorie des Masses. Fund. Math. 13 (1929), 73–116 and 333 (=Collected works, vol. 1, 599–643).

[22] A. Yu. Ol'shanskii. On the problem of the existence of an invariant mean on a group. Russian Math. Surveys 35:4 (1980), 180–181 (Russian original: Uspekhi Mat. Nauk 35:4 (1980), 199–200).

[23] A. Yu. Ol'shanskii and M. V. Sapir. Non-amenable finitely presented torsion-by-cyclic groups (submitted).